arXiv:1407.3130v2 [cs.AI] 16 Jul 2014

Allocation in Practice*

Toby Walsh

NICTA and UNSW, Sydney, Australia, email: toby.walsh@nicta.com.au

Abstract. How do we allocate scarce resources? How do we fairly allocate costs? These are two pressing challenges facing society today. I discuss two recent projects at NICTA concerning resource and cost allocation. In the first, we have been working with FoodBank Local, a social startup working in collaboration with food bank charities around the world to optimise the logistics of collecting and distributing donated food. Before we can distribute this food, we must decide how to allocate it to different charities and food kitchens. This gives rise to a fair division problem with several new dimensions, rarely considered in the literature. In the second, we have been looking at cost allocation within the distribution network of a large multinational company. This also has several new dimensions rarely considered in the literature.

1 Introduction

The next decade will throw up some fundamental and deep challenges in resource and cost allocation that computer science can help solve.

Environmental challenges: the world's resources are under increasing pressure with threats like global warning, and with the impact of an increasing population. This will require us to find ways to allocate scarce resource more efficiently and more fairly. There will also be increasing pressure to allocate costs fairly to those consuming these resources.

Economic challenges: the fall out from the global financial crisis will continue, with fresh shocks likely to occur. With growth faltering, both government and industry will increasingly focus on efficiency gains. As wealth concentrates into the hands of a few, a major and highly topical concern will be equitability.

Technological challenges: new markets enabled by the internet and mobile devices will emerge. These markets will require computational mechanisms to be developed to allocate resources and share costs fairly and efficiently.

As an example of one of these new markets, consider users sharing some resources in the cloud. How do we design a mechanism that allocates CPU and memory to those users that reflects their different preferences, and that is fair and efficient? Such a mechanism needs to be computational. We will want to implement it so that it is highly

* The work described here is joint with many colleagues including: Martin Aleksandrov, Haris Aziz, Casey Cahan, Charles Gretton, Phil Kilby, Nick Mattei. NICTA is supported by the Australian Government through the Department of Communications and the Australian Research Council through the ICT Centre of Excellence Program.

responsive and runs automatically in the cloud. And users will want to implement computational agents to bid for resources automatically. As a second example of one of these new markets, consider sharing costs in a smart grid. How do we design a mechanism that shares costs amongst users that is fair and encourages efficiency? Such a mechanism again needs to be computational. We will want to implement it so that it is highly responsive and runs automatically over the smart grid. And users will again want to implement computational agents to monitor and exploit costs in the smart grid as they change rapidly over time.

There are a number of ways in which computation can help tackle such resource and cost allocation problems. First, computation can help us set up richer, more realistic models of resource and cost allocation. Second, within such computational models, users will be able to express more complex, realistic, combinatorial preferences over allocations. Third, computation can be used to improve efficiency and equitability, and to explore the trade off between the two. And fourth, users will increasingly farm out decision making to computational agents, who will need to reason rapidly about how resources and costs are allocated.

Central to many allocation problems is a trade off between equitability and efficiency. We can, for example, give each item to the person who most values it. Whilst this is efficient, it is unlikely to be equitable. Alternatively, we can allocate items at random. Whilst this is equitable, it is unlikely to maximise utility. Rather than accept allocations that are equitable but not efficient, or efficient but not equitable, we can now use computing power to improve equitability and efficiency. Computing power can also be used to explore the Pareto frontier between equitability and efficiency.

One related area that has benefitted hugely of late from the construction of richer computational models is auctions. Billions of dollars of business have been facilitated by the development of *combinatorial* auctions, transforming several sectors including procurement and radio spectrum allocation. We expect rich new models in resource and cost allocation will drive similar transformations in other sectors. It is perhaps not surprising that a transformation has already been seen in auctions where efficiency (but not equitability) is one of the main drivers. By comparison, many of the examples we give here are in the not-for-profit and public sector where criteria like equitability are often more important. Indeed, equitability is looking increasingly likely to be a major driver of political and economic reform over the next few decades. In the not-for-profit and public sector, research like this can quickly inform both practice and policy and thereby impact on society in a major way. Indeed, such richer models are now starting to be seen in one area of resource allocation, namely kidney exchange (e.g. [1]).

2 Allocation in theory

The theoretical foundations of resource and cost allocation have been developed using simple abstract models. For example, one simple model for resource allocation is "cake cutting" in which we have a single resource that is infinitely divisible and agents with additive utility functions [2]. As a second example, a simple model for cost allocation in cooperative game theory supposes we can assign a cost to each subset of agents. As we will demonstrate shortly, abstract models like these ignore the richness and structure of

many allocation problems met in practice. For example, allocation problems are often repeated. The problem we meet today is likely to be similar to the one we will meet tomorrow. As a second example, allocation problems are often online. We must start allocating items before all the data is available. As a third example, cost functions are often complex, and dependent on time and other features of the problem. Such real world features offer both a challenge and an opportunity. For instance, by exploiting the repeated nature of an allocation problem, we may be able to increase equitability without decreasing efficiency. On the other hand, the online nature of an allocation problem makes it harder both to be efficient and to be equitable.

Our long term goal then is to develop richer, more realistic computational models for resource and cost allocation, and to design mechanisms for such models that can be fielded in practice. Many of the new applications of such models will be in distributed and asynchronous environments enabled by the internet and mobile technology. It is here that computational thinking and computational implementation is necessary, and is set to transform how we fairly and efficiently allocate costs and resources. Hence, a large focus of our work is on applying a computational lens to resource and cost allocation problems. To this end, we are concentrating on designing mechanisms that work well in practice, even in the face of fundamental limitations in the worst case.

3 Allocation in practice

We have two projects underway in NICTA which illustrate the richness of allocation problems met in practice.

3.1 Case study #1: the food bank problem

Consider the classical fair division problem. Fair division problems can be categorised along several orthogonal dimensions: divisible or indivisible goods, centralised or decentralised mechanisms, cardinal or ordinal preferences, etc [3]. Such categories are, however, not able to capture the richness of a practical resource allocation problem which came to our attention recently. FoodBank Local is a social startup founded by students from the University of New South Wales that is working with food bank charities around the world, as well as with us, to improve the efficiency of their operations. FoodBank Local won the Microsoft Imagine Cup in Australia in 2013, and were finalists worldwide for their novel and innovative approach to using technology for social good. After supermarkets, catering companies and the public have donated food, the food bank must allocate and distribute this food to charities in the local area. This requires solving a fair division problem. Fairness is important as charities often cater to different sectors of the population (geographical, social and ethnic), whilst efficiency is important to maximise the overall good. This fair division problem has several new dimensions, rarely considered in the literature.

Online: Offers of donated food arrive throughout the day. The food bank cannot wait till the end of the day before deciding how to allocate the donated food. This means we have an *online* problem, where decisions must be made before all the data is available.

Fig. 1. FoodBank Local's app for donating, allocating and distributing food.

Repeated: Each day, a food bank repeats the task of allocating, collecting and distributing food with a similar amount of donated food and set of charities. The *repeated* nature of the problem provides fresh opportunities. For example, we can be somewhat unfair today, in the expectation that we will compensate tomorrow.

Unequal entitlements: The different charities working with the food bank have different abilities at feeding their clients. The allocation of food needs to reflect this.

Combinatorial: The different charities had complex, combinatorial preferences over the donated food. A charity might want the donated apples or the bananas but not both. Models based on simple additive utilities, like those often considered in the literature, are inadequate to describe their true preferences.

Constrained: There are various constraints over the allocations. For example, we must allocate all the foods requiring refrigeration to charities served by the same truck. As a second example, certain combinations of food cannot be put together for health and safety or religious reasons.

Mixed: Each allocation problem induces a new pickup and delivery problem. This means that we have a *mixed* problem that combines resource allocation and logistics. We need to both ensure a fair division whilst at the same time optimising distribution.

To reason about such issues, we need to develop more complex and realistic models of resource allocation that borrow techniques from related areas like constraint programming [4]. Note that some of these features have individually been considered in the past. For example, we recently initiated the study of online cake cutting problems [5] in which agents being allocated the cake arrive over time. In the FoodBank problem, by comparison, it is the goods that arrive over time. Our model of online cake cutting also has none of the other features of the FoodBank problem (e.g. the FoodBank problem is repeated and preferences are combinatorial, whilst in online cake cutting, the cake is cut only once and preferences are described by a simple additive utility function). As a second example, Guo, Conitzer and Reeves have looked at repeated allocation problems [6] but their study was limited to just to one indivisible good and had none of the other features of the FoodBank problem.

We now describe our first step in building a richer model for online fair division. We stress that this is only the *first* step in putting together a richer model that has more of the features required by FoodBank Local. There are many more features which we need to add before we have a model that is close to the actual requirements. Our online model supposes items arrive over time and each item must be allocated as soon as it arrives. For simplicity, we suppose one new item arrives at each time step. Again, as a first step, we also suppose that the utility of an agent for each item is either 0 or 1. A next step would be to consider more general utilities. A simple mechanism for this model is the *like* mechanism. Agents declare whether they like an item, and we allocate each item uniformly at random between the agents that like it.

We can study the axiomatic properties of such an online fair division mechanism. For instance, it is strategy proof (the agents have an incentive to like all items with non-zero utility and no incentive to like any item with zero utility)

Proposition 1 *The like mechanism is strategy proof.*

The like mechanism is also envy free as an agent will not have greater expected utility for another agent's items.

Proposition 2 *Supposing agents bid sincerely, the like mechanism is envy free ex ante.*

However, ex post, agents can have considerable envy for the actual allocation of items given to another agent. For example, one agent could get unlucky, lose every random draw and end up being allocated no items. This agent would then envy any agent allocated item which they like. We have therefore been designing more sophisticated mechanisms which are fairer ex post. A challenge is to do so without losing a good property like strategy proofness. We also need to explore how such mechanisms work in practice, and not just in the worst case. For example, in real world problems, agents' preferences are likely to be correlated. A tool that is likely useful in such studies is identifying computational phase transitions (e.g. [7,8,9,10,11]), as well as related phenomena like backbones and backdoors (e.g. [12,13,14]).

3.2 Case study #2: cost allocation in a complex distribution network

We have come across similar rich features in real world cost allocation problems. In particular, we have been working with a large multinational company which spends hundreds of millions of dollars each year on distributing fast moving consumer goods to 20,000 customers using a fleet of 600 vehicles. They face a very challenging problem of allocating costs between customers[1]. By working with us, they have saved tens of millions of dollars per annum.

A standard method to allocate costs is to use co-operative game theory and one of the well defined cost allocation mechanisms like the Shapley value which considers the marginal cost of a customer in every possible subset of customers. In reality,

[1] These are not the actual costs charged to the customer but the costs used to decide if a customer is profitable or not. They are used as the basis for reorganising their business (e.g. changing distribution channels, renogiating contracts).

the problem is much richer than imagined in a simple abstract model like the Shapley value which supposes we can simply and easily cost each subset of customers. This richness raises several issues which are rarely considered in the literature. Several of these issues are similar to those encountered in the last example. For instance, the cost allocation problem is repeated (every day, we deliver to a similar set of customers), and constrained (since we must deliver to all supermarkets of one chain or none, we must constrain the subsets of customers to consider only those with all or none of the supermarkets in the chain). However, there are also several new issues to consider:

Computational complexity: we need a cost allocation method that is computationally easy to compute. As we argue shortly, the Shapley value is computationally challenging to compute.

Heterogeneous customers: the Shapley value supposes customers are identical, which is not the case in our problem as different customers order different amounts of product. In addition, the trucks are constrained by volume, weight, and the number of stops. We therefore need principled methods of combining the marginal cost of each unit of volume, of each unit of weight, and of each customer stop.

Complex cost functions: a cost allocation mechanism like the Shapley value supposes a simple cost function for every subset of customers. However, the cost function in our cost allocation problem is much more complex. For example, it is made up of both fixed and variable costs. As a second example, the cost function is time dependent because of traffic whilst a mechanism like the Shapley value ignores time. As a third example, the cost function depends not just on whether we deliver to a customer or not but on the channel and delivery frequency.

Strategic behaviour: cooperative game theory supposes the players are truthful and are not competing with each other. In practice, however, we have business alliances where two or more delivery companies come together to share truck space to remote areas and must then share costs. We must consider therefore that the players may game the system by misrepresenting their true goals or costs.

Sensitivity analysis: in some situations, a small change to the customer base has a very large knock-on effect on marginal costs. We are therefore interested in cost allocations that are robust to small changes in the problem.

As before, some of these features have been considered individually in the past, but combinations of these features have not. For example, [15] looked at cost allocation in a vehicle routing game with a heterogeneous fleet of vehicles. However, this study considered the solution concepts of the core and the nucleolus rather than the Shapley value. The study also ignored other features of the actual real world problem like the repeated nature of the delivery problems and the full complexity of the cost function.

A naive method to allocate costs is simply to use the marginal cost of each customer. However, such marginal costs will tend to under-estimate the actual cost. Consider two customers in the middle of the outback. Each has a small marginal cost to visit since we are already visiting the other. However, their actual cost to the business is half the cost of travelling to the outback. Cost allocation mechanism like the Shapley value deal with such problems. The Shapley value equals the average marginal cost of a customer in *every* possible subset of customers. It has nice axiomatic properties like efficiency (it allocates the whole cost), anonymity (it treats all customers alike) and monotonicity

(when the overall costs go up, no individual costs go down). However, we run into several complications when applying it to the our cost allocation problem.

One complication is that the Shapley value is computationally challenging to compute in general [16]. It involves summing an exponential number of terms (one for each possible subset of customers), and in our case each term requires solving to optimality a NP-hard routing problem. One response to the computational intractability of computing the Shapley value, is to look to approximate it. However, we have proved [17] that even finding an approximation to the Shapley value of a customer is intractable in general.

Proposition 3 *Unless P=NP, there is no polynomial time α-approximation to the Shapley value of a customer for any constant factor $\alpha > 1$.*

We have therefore considered heuristic methods based on Monte Carlo sampling [18,19], and on approximating the cost of the optimal route. We have also considered simple proxies to the Shapley value like the depot distance (that is, allocating costs proportional to the distance between customer and depot). Unfortunately we were able to identify pathological cases where there is no bound on how poorly such proxies perform. These pathologies illustrate the sort of real world features of routing problems like isolated customers which can cause difficulties.

Proposition 4 *There exists a n customer problem on which Φ_{depot}/Φ goes to 0 as n goes to ∞, and another on which Φ/Φ_{depot} goes to 0 where Φ is the true Shapley value for a customer and Φ_{depot} is the estimate for it based on the proxy of distance from depot.*

Despite such pathological problems, we were able to show that more sophisticated proxies, especially those that "blend" together several estimates, work well in practice on real world data. Figure 1 illustrates the performance of two different sampling methods. It shows that in just a few iterations we can converge on estimates that are within 5-10% of the actual Shapley value.

There are other complications that we run into when allocating costs in this domain. For example, costs are not simple but made up of fixed and variable costs which depend on the size of the delivery. As a second example, costs depend on the time of day as traffic has a large impact on the problem. As a third example, customers should not be treated equally as they order different amounts of product, and delivery vehicles are limited by both the weight and volume of product that they can deliver. As a fourth example, there are supermarket chains that require us to deliver to all their supermarkets or none. We cannot choose some profitable subset of them. All these features need ultimately to be taken into account. We are thus only at the beginning of developing richer models of cost allocation for this domain. Nevertheless, we are encouraged by these preliminary results which have resulted in significant cost savings and greatly increased profits for our business partner.

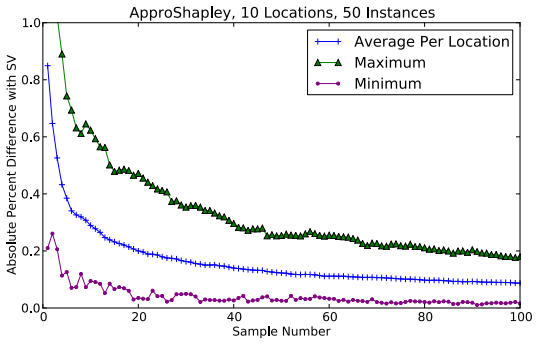

Fig. 2. Performance of the ApproShapley method [19] for Monte Carlo sampling to estimate the Shapley value in a distribution costs game. After just 100 iterations ApproShapley achieves an average error of less than 10% per location with a maximum error of less than 20%.

4 Related work

A limited number of richer models for resource and cost allocation have been considered previously. However, such models have only considered a small number of the additional features that we have described here.

One sided markets: In an one sided market, we allocate items to agents based on the preferences of the agents. The market is one sided as the items do not have preferences over the agents. It covers the allocation of goods, the division of chores, roommates and other related problems. For example, in the roommates problems, agents express preferences over each other, and the goal is to pair these agents up subject to these preferences. There has been some work already to look at more realistic models of resource allocation in one sided markets. For example, we have proposed an online version of the cake cutting problem in which agents arrive over time [5]. As a second example, [6] looked at repeated allocation problems with a single indivisible good. As a third example, [20] studied the fair division problems of indivisible goods when agents do not have completely ordered preferences over the goods, but instead have dichotomous and other succinctly specified types of preferences. As a fourth example, [21] considered roommate problems in which agents can express ties in their preference lists rather than the basic assumption of completed ordered preferences. As a fifth example, [22] set up a dynamic version of fair division, proposed some desirable axiomatic properties for such dynamic resource allocation, and designed two mechanisms that satisfy these properties. However, real world features like those met in these five examples have usually been considered in isolation. For instance, there has not been proposed a model of resource allocation in an one sided market that is simultaneously online, repeated and involves preferences that go beyond totally ordered lists.

Two sided markets: In a two sided market, both sides of the market can have preferences over each other. For example, in a stable marriage problem, the women have

preferences over the men they marry and vice versa. Another example of a two sided market is kidney exchange. There has been some work already to look at more realistic models in two sided markets. For example, in the hospital residents problem, a two sided problem in which we allocate residents to hospitals according to the preferences of the residents for the hospitals and of the hospitals for the residents, a number of real world features have been considered like couples [23], and ties in preference lists [24]. As a second example, in kidney exchange problems, a number of real world features have been considered like dynamic allocation [1], and probabilistic models of clearing (since many transplants never take place due to unforeseen logistical and medical problems) [25]. As a third example, the student-project allocation problem [26] generalises the hospital residents problem by adding capacity constraints. However, as with one side markets, such real world features have usually been considered in isolation and not in combination.

Cost allocation: Cost allocation has been widely studied in game theory and combinatorial optimisation [27,28,29]. Due to its good axiomatic properties, the Shapley value has been used in many domains [30]. A number of special cases have been identified where it is tractable to compute (e.g. [31]). However, due to the computational challenges we outlined, the Shapley value has rarely been used in the past in the sort of complex problems like the distribution problem discussed earlier. An exception is [32] which introduced the cooperative travelling salesperson (TSP) game. This also introduced the *routing game* in which the locations visited in a coalition must be traversed in a (given) fixed order. This gives a polynomial time procedure for computing cost allocations [33]. [34] developed a column generation procedure to allocate costs for a homogeneous vehicle routing problem. [15] extended this to a more practical setting of distributing gas using a *heterogeneous* fleet of vehicles. However, this study considered the solution concepts of the core and the nucleolus rather than the Shapley value. It also ignored other real world features like the repeated nature of such delivery problems. More recently [35] developed cost allocation methods for *inventory routing* problems in which customers have a capacity to hold stock, consume product at a fixed rate and the goal is to minimise costs whilst preventing stock-outs. Whilst this model has many real world features, it continues to miss others like inventory and delivery costs.

Related problems: There have been a number of complex markets developed to allocate resources which use money. For example, in a combinatorial auction, agents express prices over bundles of items [36]. Our two projects, however, only consider allocation problems where money is not transfered. Nevertheless, there are ideas from domains like combinatorial auctions which we may be able to borrow. For example, we expect the bidding languages proposed for combinatorial auction may be useful for compactly specifying complex, real world preferences even when money is not being transferred. As a second example, as occurs in some course allocation mechanisms used in practice, we can give agents virtual "money" with which to bid and thus apply an auction based mechanism [37,38].

Finally, computational phase transitions have been observed in a number of related areas including constraint satisfaction [39,40,41,42,43], number partitioning [44,45], TSP [46], social choice [47,48,49], and elsewhere [50,51,52,53,54]. We predict that a

similar analysis of phase transitions will provide insight into the precise relationship between equitability and efficiency in allocation problems.

5 Conclusions

I have discussed two recent projects at NICTA involving resource and cost allocation. Each is an allocation problem with several new dimensions, rarely considered in the literature. For example, our resource allocation problem is online, repeated, and constrained, whilst our cost allocation problem is also repeated and constrained, and additionally involves a complex cost function. These projects suggest that models for allocation problems need to be developed that are richer and more complex than the abstract models which have been used in the past to lay the theoretical foundations of the field.

References

1. Dickerson, J., Procaccia, A., Sandholm, T.: Optimizing kidney exchange with transplant chains: Theory and reality. In: Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems - Volume 2. AAMAS '12, International Foundation for Autonomous Agents and Multiagent Systems (2012) 711–718
2. Brams, S., Taylor, A.: Fair Division: From cake-cutting to dispute resolution. Cambridge University Press, Cambridge (1996)
3. Chevaleyre, Y., Dunne, P., Endriss, U., Lang, J., Lemaitre, M., Maudet, N., Padget, J., Phelps, S., Rodriguez-Aguilar, J., Sousa, P.: Issues in multiagent resource allocation. Informatica (Slovenia) **30**(1) (2006) 3–31
4. Rossi, F., van Beek, P., Walsh, T., eds.: Handbook of Constraint Programming. Foundations of Artificial Intelligence. Elsevier (2006)
5. Walsh, T.: Online cake cutting. In: Algorithmic Decision Theory, Second International Conference, ADT 2011. Proceedings. Lecture Notes in Computer Science, Springer (2011)
6. Guo, M., Conitzer, V., Reeves, D.M.: Competitive repeated allocation without payments. In: Proceedings of 5th International Workshop on Internet and Network Economics (WINE 2009). Volume 5929 of Lecture Notes in Computer Science., Springer (2009) 244–255
7. Gent, I., Walsh, T.: Phase transitions from real computational problems. In: Proceedings of the 8th International Symposium on Artificial Intelligence. (1995) 356–364
8. Gent, I., Walsh, T.: Easy problems are sometimes hard. Artificial Intelligence (1994) 335–345
9. Gent, I., Walsh, T.: The SAT phase transition. In Cohn, A.G., ed.: Proceedings of 11th ECAI, John Wiley & Sons (1994) 105–109
10. Gent, I., Walsh, T.: The hardest random SAT problems. In: Proceedings of KI-94, Saarbrucken. (1994)
11. Gent, I., Walsh, T.: The satisfiability constraint gap. Artificial Intelligence **81**(1–2) (1996)
12. Slaney, J., Walsh, T.: Backbones in optimization and approximation. In: Proceedings of 17th IJCAI, International Joint Conference on Artificial Intelligence (2001)
13. Kilby, P., Slaney, J., Thiebaux, S., Walsh, T.: Backbones and backdoors in satisfiability. In: Proceedings of the 20th National Conference on AI, Association for Advancement of Artificial Intelligence (2005)
14. Kilby, P., Slaney, J., Walsh, T.: The backbone of the travelling salesperson. In: IJCAI. (2005) 175–180

15. Engevall, S., Göthe-Lundgren, M., Värbrand, P.: The heterogeneous vehicle-routing game. Transportation Science **38**(1) (February 2004) 71–85
16. Chalkiadakis, G., Elkind, E., Wooldridge, M.: Computational aspects of cooperative game theory. Synthesis Lectures on Artificial Intelligence and Machine Learning **5**(6) (2011) 1–168
17. Aziz, H., Cahan, C., Gretton, C., Kilby, P., Mattei, N., Walsh, T.: A study of proxies for Shapley allocation of transport costs. Journal of Artificial Intelligence Research (2014) Under review.
18. Mann, I., Shapley, L.S.: Values for large games IV: Evaluating the electoral college by monte carlo. Technical report, The RAND Corporation, Santa Monica, CA, USA (1960)
19. Castro, J., Gomez, D., Tejada, J.: Polynomial calculation of the Shapley value based on sampling. Computers & Operations Research **36**(5) (2009) 1726 – 1730
20. Bouveret, S., Lang, J.: Efficiency and envy-freeness in fair division of indivisible goods: Logical representation and complexity. Journal of Artificial Intelligence Research (JAIR) **32** (2008) 525–564
21. Irving, R., Manlove, D.: The stable roommates problem with ties. J. Algorithms **43**(1) (2002) 85–105
22. Kash, I., Procaccia, A., Shah, N.: No agent left behind: Dynamic fair division of multiple resources. In: Proceedings of the 2013 International Conference on Autonomous Agents and Multi-agent Systems. AAMAS '13, Richland, SC, International Foundation for Autonomous Agents and Multiagent Systems (2013) 351–358
23. McDermid, E., Manlove, D.: Keeping partners together: algorithmic results for the hospitals/residents problem with couples. Journal of Combinatorial Optimization **19**(3) (2010) 279–303
24. Irving, R., Manlove, D., Scott, S.: The hospitals/residents problem with ties. In: Algorithm Theory - SWAT 2000. Volume 1851 of Lecture Notes in Computer Science. Springer Berlin Heidelberg (2000) 259–271
25. Dickerson, J., Procaccia, A., Sandholm, T.: Failure-aware kidney exchange. In: Proceedings of the Fourteenth ACM Conference on Electronic Commerce. EC '13, New York, NY, USA, ACM (2013) 323–340
26. Abraham, D., Irving, R., Manlove, D.: Two algorithms for the student-project allocation problem. Journal of Discrete Algorithms **5**(1) (2007) 73 – 90
27. Koster, M.: Cost sharing. Encyclopedia of Complexity and Systems Science (2009) 724–753
28. Curiel, I.: Cooperative combinatorial games. In Chinchuluun, A., Pardalos, P., Migdalas, A., Pitsoulis, L., eds.: Pareto Optimality, Game Theory And Equilibria. Volume 17 of Springer Optimization and Its Applications. Springer New York (2008) 131–157
29. Young, H.P.: Producer incentives in cost allocation. **53**(4) (1985) 757–765
30. Young, H.P.: Cost allocation. In: Handbook of Game Theory with Economic Applications. Volume 2. (1994) 1193–1235
31. Mann, I., Shapley, L.S.: Values for large games IV: Evaluating the electoral college exactly. Technical report, The RAND Corporation, Santa Monica, CA, USA (1962)
32. Potters, J.A., Curiel, I.J., Tijs, S.H.: Traveling salesman games. Mathematical Programming **53**(1-3) (1992) 199–211
33. Derks, J., Kuipers, J.: On the core of routing games. International Journal of Game Theory **26**(2) (1997) 193–205
34. Göthe-Lundgren, M., Jörnsten, K., Värbrand, P.: On the nucleolus of the basic vehicle routing game. Mathematical Programming **72**(1) (1996) 83–100
35. Özener, O.O., Ergun, O., Savelsbergh, M.: Allocating cost of service to customers in inventory routing. Oper. Res. **61**(1) (January 2013) 112–125
36. Cramton, P., Shoham, Y., Steinberg, R.: Combinatorial Auctions. MIT Press (2006)

37. Sonmez, T., Unver, M.: Course bidding at business schools. International Economic Review **51**(1) (2010) 99–123
38. Budish, E., Cantillon, E.: The Multi-Unit Assignment Problem: Theory and Evidence from Course Allocation at Harvard. American Economic Review **102**(5) (2012) 2237–71
39. Gent, I., MacIntyre, E., Prosser, P., Walsh, T.: Scaling effects in the CSP phase transition. In: 1st International Conference on Principles and Practices of Constraint Programming (CP-95), Springer-Verlag (1995) 70–87
40. Gent, I., MacIntyre, E., Prosser, P., Walsh, T.: The constrainedness of search. In: Proceedings of the 13th National Conference on AI, Association for Advancement of Artificial Intelligence (1996) 246–252
41. Gent, I., MacIntyre, E., Prosser, P., Walsh, T.: The scaling of search cost. In: Proceedings of the 14th National Conference on AI, Association for Advancement of Artificial Intelligence (1997) 315–320
42. Walsh, T.: The constrainedness knife-edge. In: Proceedings of the 15th National Conference on AI, Association for Advancement of Artificial Intelligence (1998)
43. Gent, I., MacIntyre, E., Prosser, P., Smith, B., Walsh, T.: Random constraint satisfaction: Flaws and structure. Constraints **6**(4) (2001) 345–372
44. Gent, I., Walsh, T.: Phase transitions and annealed theories: Number partitioning as a case study. In: Proceedings of 12th ECAI. (1996)
45. Gent, I., Walsh, T.: Analysis of heuristics for number partitioning. Computational Intelligence **14**(3) (1998) 430–451
46. Gent, I., Walsh, T.: The TSP phase transition. Artificial Intelligence **88** (1996) 349–358
47. Walsh, T.: Where are the really hard manipulation problems? The phase transition in manipulating the veto rule. In: Proceedings of 21st IJCAI, International Joint Conference on Artificial Intelligence (2009) 324–329
48. Walsh, T.: An empirical study of the manipulability of single transferable voting. In Coelho, H., Studer, R., Wooldridge, M., eds.: Proc. of the 19th European Conference on Artificial Intelligence (ECAI-2010). Volume 215 of Frontiers in Artificial Intelligence and Applications., IOS Press (2010) 257–262
49. Walsh, T.: Where are the hard manipulation problems? Journal of Artificial Intelligence Research **42** (2011) 1–39
50. Walsh, T.: Search in a small world. In: Proceedings of 16th IJCAI, International Joint Conference on Artificial Intelligence (1999)
51. Gent, I., Walsh, T.: Beyond NP: the QSAT phase transition. In: Proceedings of the 16th National Conference on AI, Association for Advancement of Artificial Intelligence (1999)
52. Gent, I., Hoos, H., Prosser, P., Walsh, T.: Morphing: Combining structure and randomness. In: Proceedings of the 16th National Conference on AI, Association for Advancement of Artificial Intelligence (1999)
53. Walsh, T.: Search on high degree graphs. In: Proceedings of 17th IJCAI, International Joint Conference on Artificial Intelligence (2001)
54. Walsh, T.: From P to NP: COL, XOR, NAE, 1-in-k, and Horn SAT. In: Proceedings of the 17th National Conference on AI, Association for Advancement of Artificial Intelligence (2002)